Workshop on

Mathematics and Machines: Navigating the Transition from Brown to Sustainable Finance

Book of Abstracts

Lorentz Center · Leiden, The Netherlands · 26–29 May 2026

Contents

1 Venue and sponsors

The workshop is hosted and supported by the Lorentz Center.

We are grateful to the sponsors for their financial support of the workshop.

2 Program

Tuesday, May 26, 2026 – Algorithmic Trading

10:00–10:15	Opening session	
10:15–11:45	**Johannes Muhle-Karbe**	*A Unified Theory of Order Flow, Market Impact, and Volatility I*
11:45–12:15	**Mathieu Rosenbaum**	*A Unified Theory of Order Flow, Market Impact, and Volatility II*
12:15–13:15	Lunch Break	
13:15–13:45	**Rob Graumans**	*AI and the Challenge of Supervising Algorithmic Trading*
13:45–14:30	**Panel discussion**	*Algorithmic Trading: from Theory to Practice*
	Panelists:	Rob Graumans *(AFM)* Rik Lodder *(MN)* Johannes Muhle-Karbe *(Imperial College)* Martin van der Schans *(Robeco)*
	Moderator:	Antonis Papapantoleon *(TU Delft)*
14:30–14:45	Break	
14:45–15:15	**Eduardo Abi Jaber**	*Fading Memory Signatures for Sequential Modeling*
15:15–15:45	**Christa Cuchiero**	*Dynamic universal approximation and modeling with neural and signature SDEs*
15:45–16:15	Break	
16:15–17:00	**Group Work**	*Parallel sessions*
	Session A:	Discussion on collaboration and funding opportunities
	Session B:	Short talks (Janusz Meylahn & Fenghui Yu)
17:00–18:30	**Welcome reception**	*Poster session*

Wednesday, May 27, 2026 – Sustainable and Green Finance

09:00–10:30	**René Aïd**	*Information Provision in Energy and Carbon Finance*
10:30–11:00	Break	
11:00–11:30	**Giorgia Callegaro**	*Continuous-time persuasion by filtering*
11:30–12:00	**Koos Gubbels**	*Financial impacts of extreme weather from an actuarial perspective*
12:00–13:15	Lunch Break	
13:15–14:15	**Panel discussion**	*Climate Risk Across the Spectrum*
	Panelists:	Hielke van der Aa *(NWB)* Gijs Kloek *(Achmea)* Martine Koridon-Rommen *(MN)* Bert Kramer *(ORTEC)*
	Moderator:	Zac Taylor *(TU Delft)*
14:15–14:45	**Peter Tankov**	*Debt-for-Climate Swaps and Adaptation Investment in a Stochastic Sovereign Debt Model*
14:45–15:15	**Olivier Guéant**	*Buy it, Store it, Sell it: On the Optimality Gap of the Rolling Intrinsic Strategy*
15:15–15:45	Break	
15:45–16:30	**Group Work**	*Parallel sessions*
	Session A:	Discussion on collaboration and funding opportunities
	Session B:	Short talks (Sven Karbach, Alexandros Saplaouras & Özge Şahin)
16:30–18:00	**Practitioner's afternoon**	Charles Martinez *(G-Research)* Daniel Giamouridis *(Qube RT)* Kin Lee *(MN)* NN *(Cross Options)* Drona Kandhai *(ING)*

Thursday, May 28, 2026 – Machine Learning and Stochastic Control in Finance

09:00–10:30	**Blanka Horváth**	*Generative Models and Operator Learning for Implied Volatility Surfaces*
10:30–11:00	Break	
11:00–11:30	**Christian Bayer**	*Global and local regression on path space: a signature approach with applications*
11:30–12:00	**Iman Honarvar**	*The State of Machine Learning and AI in Modern Asset Management*
12:00–13:10	Lunch Break	
13:10–13:15	**Kees Vuik**	Presentation of ICIAM 2027
13:15–14:15	**Panel discussion**	*Machine Learning in Finance: Hype and Reality*
	Panelists:	Sjors Altemuhl *(DNB)*
		Fang Fang *(FF Quant & TU Delft)*
		Diederik Fokkema *(ADC)*
		Roger Lord *(Cardano)*
	Moderator:	Peter Forsyth *(Waterloo)*
14:15–14:45	**Geneviève Gauthier**	*Deep Implied Volatility Factor Models for Stock Options*
14:45–15:15	**Irene Monasterolo**	*Climate credit risk and corporate valuation: insights from the Network for Greening the Financial System short-term climate scenarios*
15:15–15:45	Break	
15:45–16:15	**Group Work**	*Parallel sessions*
	Session A:	Discussion on collaboration and funding opportunities
	Session B:	Short talks (Costas Smaragdakis & Shuaiqiang Liu)
	Session C:	Short talks (Nestor Parolya, Laura Spierdijk)
16:15–16:45	**Dylan Possamaï**	*Contract theory and energy markets*
17:00–	**Excursion & Dinner**	

Friday, May 29, 2026 – FinTech: Modeling and Trading in Decentralized Exchanges

09:00–10:30	**Emmanuel Gobet**	*Quantitative Modelling of Risks in Decentralized Finance*
10:30–11:00	Break	
11:00–11:30	**Natalie Packham**	*Sentiment-driven forecasting of fractional trading activity using LSTM networks*
11:30–12:00	**Jasper Anderluh**	*Investing by and for Future Generations*
12:00–13:00	Lunch Break	
13:00–13:30	**Yuying Li**	*Optimal Stochastic Dynamic Strategies: an Across-Time-Neural-Network Approach*
13:30–14:00	**Athena Picarelli**	*Extended mean field control: a global numerical solution via finite-dimensional approximation*
14:00–15:00	**Group Work** *Session A:* *Session B:*	*Parallel sessions & Wrap-up* Discussion on collaboration and funding opportunities Short talks (Evgenii Vladimirov & Michael Samet)

3 Mini-courses

René Aid (Universite Paris Dauphine)

Information Provision in Energy and Carbon Finance

The objective of this tutorial is to introduce the regulatory and economic issues of carbon finance and energy markets. The tutorial first reviews the main regulatory frameworks and policy challenges shaping these markets, and in particular, the role of information. Second, it presents some models showing how information provision can shape agents decisions of decarbonization.

Johannes Muhle-Karbe (Imperial College London)

A Unified Theory of Order Flow, Market Impact, and Volatility I

We propose a microstructural model for the order flow in financial markets that distinguishes between core orders and reaction flow, both modeled as Hawkes processes. This model has a natural scaling limit that reconciles a number of salient empirical properties: persistent signed order flow, rough trading volume and volatility, and power-law market impact. In our framework, all these quantities are pinned down by a single statistic H_0, which measures the persistence of the core flow. Specifically, the signed flow converges to the sum of a fractional process with Hurst index H_0 and a martingale, while the limiting traded volume is a rough process with Hurst index $H_0 - 1/2$. No-arbitrage constraints imply that volatility is rough, with Hurst parameter $2H_0 - 3/2$, and that the price impact of trades follows a power law with exponent $2-2H_0$. The analysis of signed order flow data yields an estimate H_0 around 3/4. This is not only consistent with the square-root law of market impact, but also turns out to match estimates for the roughness of traded volumes and volatilities remarkably well.

(Based on joint work with Youssef Ouazzani Chahdi, Mathieu Rosenbaum and Grégoire Szymanski available at `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6155066`.)

Blanka Horvath (University of Oxford)

Generative Models and Operator Learning for Implied Volatility Surfaces

Modern financial markets generate complex, high-dimensional data streams whose structure differs fundamentally from the sequential data encountered in other machine learning domains. This mini-course explores the interplay between generative modelling, operator learning, and stochastic analysis in the context of financial market dynamics, with a particular focus on the implied volatility surface. We discuss how classical calibration and interpolation problems can be recast as learning problems over function spaces, and examine what structural properties — such as no-arbitrage constraints and path-dependence — must be built into or enforced by such learners. We highlight connections between deep learning approaches and classical mathematical finance, and discuss directions for building principled, scalable, and financially meaningful models for surface dynamics and market generation.

Emmanuel Gobet (Sorbonne Université)

Quantitative modelling of risks in Decentralized Finance

This minicourse provides an overview of the quantitative aspects of decentralized finance (DeFi). After a brief introduction to blockchain technology and the smart contracts that underpin it, the course is organized around three main topics: (i) data description and aggregation, including classification and outlier detection; (ii) automated market makers, with a focus on general principles and selected results on Uniswap; and (iii) lending and borrowing protocols, with emphasis on general mechanisms and selected results on Aave. For each topic, we present relevant modeling approaches and discuss the theoretical and practical challenges they raise.

4 Invited talks

Eduardo Abi Jaber (École Polytechnique)

Fading Memory Signatures for Sequential Modeling

We introduce the Exponentially Fading Memory (EFM) Signature, a novel tool for sequential modeling and learning in time series. Its versatility is demonstrated across applications such as electricity demand forecasting, modeling VIX, and regime change detection, where it captures both short and long-term dependencies and improves predictive performance. We construct the EFM-signature via rough path theory, carefully adapted to accommodate improper integration from minus infinity. The EFM-signature retains many of the key algebraic and analytical properties of classical signatures, including a suitably modified Chen identity, the linearization property, path-determinacy, and the universal approximation property. From the probabilistic perspective, the EFM-signature provides a "stationarized" representation, making it particularly well-suited for timeseries analysis and signal processing overcoming the shortcomings of the standard signature. In particular, the EFM-signature of time-augmented Brownian motion evolves as a group-valued Ornstein–Uhlenbeck process. We establish its stationarity, Markov property, and exponential ergodicity in the Wasserstein distance, and we derive an explicit formula a la Fawcett for its expected value in terms of Magnus expansions. Based on joint works with Francesca Primavera and Dimitri Sotnikov.

Christian Bayer (Weierstrass Institute, Berlin)

Global and local regression on path space: a signature approach with applications

The path signature is a powerful tool for solving regression problems on path space, i.e., for computing conditional expectations $E[Y|X]$ when the random variable X is a stochastic process – or a time-series. We provide new theoretical convergence guarantees for two different, complementary approaches to regression using signature methods. In the context of global regression, we show that linear functionals of the robust signature are universal in the L^p sense in a wide class of examples. In addition, we present a local regression method based on signature semi-metrics, and show universality as well as rates of convergence.

Giorgia Callegaro (University of Padova)

Continuous-time persuasion by filtering

We frame dynamic persuasion in a partial observation stochastic control Leader-Follower game with an ergodic criterion. The Receiver controls the dynamics of a multidimensional unobserved state process. Information is provided to the Receiver through a device designed by the Sender that generates the observation process. The commitment of the Sender is enforced. We develop this approach in the case where all dynamics are linear and the preferences of the Receiver are linear-quadratic. We prove a verification theorem for the existence and uniqueness of the solution of the HJB equation satisfied by the Receiver's value function. An extension to the case of persuasion of a mean field of interacting Receivers is also provided. We illustrate this approach in two applications: the provision of information to electricity consumers with a smart meter designed by an electricity producer; the information provided by carbon footprint accounting rules to companies engaged in a best-in-class emissions reduction effort. In the first application, we link the benefits of information provision to the mispricing of electricity production. In the latter, we show that even in the absence of information cost, it might be optimal for the regulator to blur information available to firms to prevent them from coordinating on a higher level of carbon footprint to reduce their cost of reaching a below average emission target.

Christa Cuchiero (University of Vienna)

Dynamic universal approximation and modeling with neural and signature SDEs

Motivated by generative AI inspired modeling, we investigate the mathematical foundations of neural and signature stochastic differential equations (SDEs). In these models, the coefficients of the SDE are parameterized either by neural networks or by functions of the path signature, leveraging their well-known universal approximation properties. While such classical universal approximation results concern static functions, our focus is on dynamic universal approximation at the level of the SDE solutions. We show that any standard SDE as well as path-dependent SDE can be approximated by a suitable neural or signature SDE, respectively. For signature SDEs, this requires the development of a dedicated well-posedness theory. Building on this foundation, we prove in particular that SDEs whose coefficients are entire functions of the signature serve as dynamic universal approximators for path-dependent SDEs. Since these signature SDEs exhibit an affine structure with respect to the lifted signature state, this result implies universality of affine processes within the class of Itô-processes. We also outline how a similar approach can be used to tackle generic path-dependent optimal control problems

Geneviève Gauthier (HEC Montréal)

Deep Implied Volatility Factor Models for Stock Options

This paper proposes the Deep Implied Volatility Factor Model to estimate daily implied volatility surfaces for individual stocks. The model combines neural networks to learn common latent functions with linear regression for time-varying factor loadings, preserving interpretability. Incorporating time-to-earnings-announcement information captures event-driven volatility dynamics and associated shifts in the risk-neutral density. The framework enables consistent derivative pricing and the construction of firm-level VIX-style indices that closely track official measures and extend volatility estimation to periods with sparse option data. Evidence from Apple and five additional stocks demonstrates robustness and scalability.

Olivier Guéant (LPSM, Université Paris Cité)

Buy it, Store it, Sell it: On the Optimality Gap of the Rolling Intrinsic Strategy

The growing penetration of renewable energy sources is associated with significant volatility in electricity prices, particularly on intraday markets. In this evolving environment, battery energy storage systems (BESS) have emerged as key assets, enabling operators to generate profit by buying low, storing energy, and selling high, while contributing to price smoothing. Among the various operational strategies used to determine charging and discharging schedules, the rolling intrinsic strategy is widely adopted in practice due to its simplicity and intuitive appeal. In this paper, we develop a quantitative framework to evaluate the performance of the rolling intrinsic strategy relative to optimal strategies derived from stochastic control formulations under varying degrees of risk aversion. Our results provide new insights into the optimality gap associated with the rolling intrinsic approach and highlight the importance of model risk in the design of the decision-making process of BESS operators.

Yuying Li (University of Waterloo)

Optimal Stochastic Dynamic Strategies: an Across-Time-Neural-Network Approach

We present a data-driven neural network framework for computing optimal stochastic investment strategies over long horizons. Our approach applies to diverse financial optimal dynamic decision problems. Our across-time neural network architecture learns optimal dynamic policies directly from data, avoiding computationally expensive dynamic programming methods that require high-dimensional conditional expectations. Instead, we solve a low-dimensional control problem through single optimization. Using inflation-adjusted CRSP market data, we demonstrate the performance and practical relevance of the optimal strategies in trading leveraged ETFs in a long horizon investment portfolio as well as in DC pension decumulation. Note. This talk is based on joint work with P. Forsyth (Univ of Waterloo), M. Chen (JP Morgan), M. Shirazi (Bank of Canada), and P. Van Staden (National Australia Bank, Melbourne)

Irene Monasterolo (Utrecht University)

Climate credit risk and corporate valuation: insights from the Network for Greening the Financial System short-term climate scenarios

We develop a micro-founded model of climate credit risk (CLIMACRED) for scenario-contingent valuation, linking adjustments of the firm's default probability and credit risk to the reference supervisory climate transition scenarios of the Network for Greening the Financial System (NGFS). We unveil a new economic mechanism of asset stranding, with transition risk emerging from the change in asset value triggered by a change in markets' expectations about the material impact of a climate policy scenario on the firm. We quantify the impact of changes in market's expectations about climate policy scenarios on the financial valuation adjustment of the firm's financial instruments (equity and bonds), applying CLIMACRED to a universe of utility companies. Then, we present the application of the CLIMACRED model for the development of the NGFS short-term (ST) climate scenarios that quantify the impact of extreme weather events, transition risk, and their interplay, on 46 countries, and 50 economic sectors. The short-term horizon, the granular assessment of acute and compound risk, and the scenario-contingent adjustment of financial valuation of securities and contracts, make the scenarios more aligned with stress-testing exercises

Natalie Packham (HWR Berlin)

Sentiment-driven forecasting of fractional trading activity using LSTM networks

Fractional trading has emerged as an opportunity for investors to invest in fractions of shares, enabling access to high-priced stocks even with small budgets, and thus reaching a bigger investment community. Aside from the inventory risk faced by all market makers, providers of fractional trading are unable to perfectly hedge their positions and therefore need to forecast retail investors' order flow over short time horizons to manage their market risk. We develop machine learning models to predict trading activity in fractional trading using time series data on investor holdings of more than 8,000 securities combined with daily news sentiment indicators. Our approach proceeds in two stages. First, secu- rities are grouped into clusters with similar trading patterns, and representative firms are selected from each cluster. Second, trading activity is modeled using long short-term memory (LSTM) neural networks augmented with sentiment features. The results show that fractional trading activity can be modeled with meaningful predictive accuracy and that incorporating news sentiment improves LSTM performance in most cases. Feature importance analysis using methods from explainable artificial intelligence (XAI) identifies the primary drivers of trading activity.

Athena Picarelli (University of Verona)

Extended mean field control: a global numerical solution via finite-dimensional approximation

We investigate the global numerical approximation of a class of extended mean field control problems (MFC), where the dynamics and costs depend on the joint distribution of the state and the control. We propose a framework to approximate the value function globally over the Wasserstein space, moving beyond the restriction of fixed initial conditions. Our approach exploits the propagation of chaos by approximating the infinite-dimensional MFC problem by an N-player cooperative game, together with the usage of finite-dimensional solvers. This method avoids the need to parametrise functions on an infinite-dimensional space, offering a balance between probabilistic rigor and computational efficiency.

Dylan Possamaï (ETH Zürich)

Contract theory and energy markets

We study how continuous-time contract theory can be used as a quantitative tool for the design of incentives in energy markets, focusing on two applications: demand–response management and pollution regulation. First, we model a demand–response program as a moral-hazard Principal–Agent problem in which a risk-averse consumer can exert costly effort to reduce both the mean and the volatility of her consumption deviations, while the producer observes consumption but not effort. We then extend the model to a large population with common shocks, leading to a mean-field Stackelberg setting where optimal incentives naturally combine (i) a "classical" component indexed on the individual consumption deviation and (ii) a benchmarking component indexed on aggregate behavior (equivalently, on the common noise). This yields explicit contract structures and comparative statics, and illustrates how controlling volatility can materially increase responsiveness. Second, we propose a regulator–producers model for pollution regulation in an electricity network. After an auction phase (cost functions given), the regulator chooses production and power flows subject to network constraints, and designs terminal transfers that induce producers to exert abatement effort affecting the drift of aggregate emissions. The set of implementable remunerations is characterized via a tractable dynamic representation, while the regulator's problem leads to an HJB characterization and computable optimal contracts. Numerical experiments on a calibrated three-node network (inspired by the Chilean market) show that optimal contracts can substantially curb emissions, reducing pollution increments by more than 30% in the reported simulation study.

Mathieu Rosenbaum (Université Paris Dauphine)

A unified theory of order flow, market impact and volatility, II

We propose a microstructural model for the order flow in financial markets that distinguishes between core orders and reaction flow, both modeled as Hawkes processes. This model has a natural scaling limit that reconciles a number of salient empirical properties: persistent signed order flow, rough trading volume and volatility, and power-law market impact. In our framework, all these quantities are pinned down by a single statistic H_0, which measures the persistence of the core flow. Specifically, the signed flow converges to the sum of a fractional process with Hurst index H_0 and a martingale, while the limiting traded volume is a rough process with Hurst index $H_0 - 1/2$. No-arbitrage constraints imply that volatility is rough, with Hurst parameter $2H_0 - 3/2$, and that the price impact of trades follows a power law with exponent $2 - 2H_0$. The analysis of signed order flow data yields an estimate H_0 close to $3/4$. This is not only consistent with the square-root law of market impact, but also turns out to match estimates for the roughness of traded volumes and volatilities remarkably well. This is joint work with Johannes Muhle Karbe, Youssef Ouazzani-Chahdi, and Gregoire Szymanski.

Peter Tankov (ENSAE, Institut Polytechnique de Paris)

Debt-for-Climate Swaps and Adaptation Investment in a Stochastic Sovereign Debt Model

Debt-for-climate swaps are financial arrangements in which creditors grant debt relief to a sovereign in exchange for a commitment to finance climate adaptation or resilience investment. We propose a dynamic stochastic framework for analyzing such instruments in small climate-vulnerable economies. In the model, climate shocks raise public debt through reconstruction needs and reduce future output, while sovereign credit spreads and default intensity depend endogenously on debt sustainability. The government chooses spending and adaptation investment, and we compare direct resilience financing with debt-for-climate swaps that relax fiscal constraints by creating borrowing space for adaptation. The analysis highlights a fiscal-space channel: even when resilience investment has high long-run value, it may be postponed when debt is already high, whereas debt relief can bring investment forward and improve debt dynamics at lower cost to creditors than a full grant.

5 Industry talks

Jasper Anderluh (TU Delft, Woolsocks, Frank Energy)

Investing by and for Future Generations

Nowadays, young people start earlier with investing, i.e. 30% of Gen Z begin investing in university or early adulthood, compared to 15% of Millennials and 6% of Baby Boomers. Of those young investors, over 80% are interested in parental guidance. This talk focuses on the dynamic mix of technical possibilities, increased investment regulation and different needs for (parentally guided) investments of the future generations.

Rob Graumans (AFM)

AI and the Challenge of Supervising Algorithmic Trading

This talk discusses how recent advances in AI are changing the supervision of algorithmic trading. Drawing on joint work with academic researchers, it also presents new evidence on "flickering quotes" — orders with extremely short resting times — showing that, despite their fleeting nature, they can have a lasting impact on prices. These findings highlight the need to rethink how supervisors detect and assess market behavior in increasingly fast and complex trading environments.

Koos Gubbels (Achmea)

Financial impacts of extreme weather from an actuarial perspective

We consider the financial impacts of extreme weather from an actuarial perspective. Extreme weather events, such as wind storms, cloud bursts and hail storms, can cause large financial losses, most of which are covered by insurance companies. As a result, insurers have to price, reserve and hold capital for these extreme events, whose frequency and intensity are increasing. In recent years, Achmea has started a research programme to combine insurance data with meteorological data to assess the financial impacts of extreme weather. We discuss the insights from this research programme for actuarial modelling of extreme weather claims, including micro-level reserving models and natural catastrophe capital models.

Iman Honarvar (Robeco)

The State of Machine Learning and AI in Modern Asset Management

This talk reviews the current state of machine learning and artificial intelligence in asset management from an industry perspective. I discuss practical applications of ML, NLP, and generative AI across return prediction, risk management, sustainability investing, and trading, highlighting where these methods meaningfully add value and where they fall short. Emphasis is placed on data limitations, non-stationarity, and the complementarity between human judgment and machine intelligence. The goal is to connect academic advances with real-world investment practice.

6 Short talks

Sven Karbach (University of Amsterdam)

Semi-Static hedging of volumetric risk in energy markets

In this talk, we develop quantitative methods for pricing and hedging Power Purchase Agreements (PPAs), whose value depends on the joint dynamics of future renewable production and forward electricity prices. We propose a coupled HJM framework for forward power prices and renewable production indices, driven by a Wishart-type stochastic covariance model that captures their complex dependence structure. Within this setting, we derive semi-closed-form solutions for the variance-optimal hedge and assess its effectiveness in mitigating intrinsic volume and price risks. Our integrated approach is benchmarked against Delta hedging and a fully static strategy based on portfolios of power and weather derivatives.

Shuaiqiang Liu (Delft University of Technology & ING Bank, the Netherlands)

Controllable Generation of Implied Volatility Surfaces with Deep Generative models

Implied volatility surfaces (IVSs) are a critical input to a wide range of financial and managerial decisions, including derivative pricing, hedging, risk aggregation, and regulatory stress testing. These decisions are most consequential during periods of market stress, when IVSs exhibit extreme behavior and historical data are scarce. Existing data-driven and model-based approaches often do not perform well in such regimes, providing limited control over economically meaningful surface characteristics under regulatory constraints. We present a deep generative modelling framework for controllably synthesizing IVSs that explicitly support financial and regulatory decision making under stress. Building on variational autoencoders (VAE), we design a controllable VAE architecture, in which meaningful IVS shape characteristics are treated as interpretable control variables, while latent variables capture residual structure to preserve shape diversity. To enable this control, IVS feature values are quantified at a selected anchor point and incorporated into the decoder to steer generation. The framework ensures that generated IVSs remain economically consistent and suitable for downstream pricing and risk management applications. Using S&P 500 option data and a COVID-19 stress-period case study, we show that the proposed approach accurately controls prescribed surface features, and captures market-specific volatility patterns that are absent from simulated data. Our results demonstrate that controllable generative AI models can serve as a practical decision-support tool for risk managers and regulators.

Janusz Meylahn (University of Twente)

Algorithmic Collusion in Finance

Recent work on multiagent learning asks whether it is possible for algorithms to learn to collude in pricing environments. Some results suggest that this may be possible. Given the prevalence of algorithms in financial markets, this raises the question of whether this phenomenon is possible there. In this talk I briefly discuss previous work in pricing environments and then outline what the main challenges are likely to be when moving to financial markets.

Özge Şahin (TU Delft)

From hundreds to dozens: Stochastic hill climbing for parsimonious Environmental, Social, and Governance scoring

Environmental, Social, and Governance (ESG) metrics are widely adopted tools for assessing corporate sustainability. Despite their importance for investors, regulators, and other stakeholders, ESG scores suffer from major drawbacks, including score divergence across providers, opaque and biased methodologies, and data quality issues. This paper focuses on Eikon Refinitiv, a prominent ESG data provider, to explore two interrelated challenges: the distortion of association measures due to missing data imputation and the potential redundancy among ESG Key Performance Indicators (KPIs). First, we show that Refinitiv's percentile ranking scheme preserves rank correlations (Kendall's tau, Spearman's rho) but not Pearson's correlation when data are complete, but Refinitiv's imputation for missing values with zeros inflates all association measures. Second, we develop an optimization-based tool to identify informational redundancy among KPIs. We formulate it as a cardinality-constrained rank correlation maximization problem and solve it via stochastic hill climbing with random restarts. Across three sectors and ten years, we find that only 20–25% of KPIs are needed to approximate pillar scores, suggesting that ESG scoring models can be substantially streamlined without sacrificing accuracy. Our theoretical and empirical findings have important implications for regulatory efforts to standardize ESG ratings and offer practical guidance for constructing interpretable, reliable, and efficient ESG scores.

Nestor Parolya (Delft University of Technology)

Pseudo-inverses as regularizers in high-dimensional learning for finance

Many machine learning and data-driven optimization methods in finance operate in regimes where the number of variables is comparable to or larger than the sample size, making covariance estimation and matrix inversion unstable. In this talk, I show that the Moore–Penrose inverse and related ridge-type inverses admit a useful high-dimensional interpretation as regularized versions of covariance inversion. Recent asymptotic results provide explicit deterministic equivalents for weighted trace functionals under weak distributional assumptions and general population covariance structure. These findings lead to simple fully data-driven shrinkage rules and help explain why pseudo-inverse methods can improve stability in practice. I will illustrate the main idea with an application to large minimum variance portfolio estimation, where the resulting procedures yield more stable allocations and improved out-of-sample performance.

Michael Samet (RWTH Aachen)

Data-Driven Stochastic Optimal Control for Intraday Electricity Trading by Renewable Producers

The rapid growth of weather-dependent renewable generation increases price volatility and imbalance penalty risk in power markets, creating the need for advanced quantitative trading strategies. We develop a data-driven continuous-time stochastic optimal control framework for intraday electricity trading using stochastic differential equations with drift terms ensuring mean reversion to deterministic forecast trajectories. Production follows a Jacobi diffusion, while prices follow an asymmetric jump-diffusion to reflect the heavy-tailed behavior observed in intraday markets. The framework accounts for realistic market features by incorporating gate closure and energy-based imbalance settlement over the delivery window, where the path-dependent imbalance cost is handled by state augmentation to preserve the Markovian structure. The value function is characterized via the dynamic programming principle by a three-stage sequence of two linear Kolmogorov backward equations and a nonlinear Hamilton-Jacobi-Bellman partial integro-differential equation. To solve this problem efficiently, we propose a monotone IMEX finite-difference scheme with operator splitting, semi-implicit linearization, and a differential formulation for the jump operator. Numerical experiments based on German market data show that the optimal trading strategy significantly outperforms the benchmark time-weighted average price strategy and achieves performance close to the perfect-foresight benchmark. Sensitivity experiments further show how jump intensity, delivery-window length, and trading horizon affect the trading policy and the resulting profit-and-loss distribution.

Alexandros Saplaouras (University of the Aegean)

Predicting energy production using SDEs via classical methods and machine learning

We propose a short-term forecasting framework for wind power production that combines stochastic differential equations with machine learning. Starting from Numerical Weather Predictions, we introduce a mean-reverting stochastic correction model whose parameters are efficiently estimated from SCADA data using an Euler-based scheme inspired by classical CIR dynamics. Power output is obtained via direction-dependent power curves refined using either neural networks or classical interpolation methods. The proposed approach highlights the role of stochastic analysis in addressing challenges in energy markets and sustainable finance.

Costas Smaragdakis (University of the Aegean & FORTH)

Runge–Kutta Physics-Informed Neural Networks for evolutionary PDEs

We introduce a Runge–Kutta Physics-Informed Neural Network framework designed for the stable and accurate solution of evolutionary PDEs. Standard neural network solvers often struggle with stiff dynamics or fail to preserve conservation laws. To address this, we propose a new approach that combines collocation Runge–Kutta time discretization with a continuous polynomial reconstruction of the neural network predictions. By enforcing the PDE dynamics strictly at specific quadrature nodes (Gauss, Radau IIA, or Lobatto IIIA), the proposed method ensures that our solver inherits the stability properties of the underlying implicit Runge–Kutta schemes.

Laura Spierdijk (University of Twente)

Characterizing and Computing Efficient Portfolios: A Stochastic Dominance Approach

Stochastic dominance (SD) provides a framework for ranking uncertain outcomes in a manner consistent with expected utility theory. However, the absence of practical computational methods to identify the nth-order SD-efficient set has limited its widespread application in portfolio optimization. This study fills this gap by providing a complete mathematical characterization of the SD-efficient set for orders two and higher; that is, the set of all portfolios not dominated at that order. We derive key topological properties of this set and use them to construct a neighborhood-based search algorithm that approximates the SD-efficient set. By focusing primarily on efficient portfolios, the algorithm achieves substantial computational gains compared with exhaustive pairwise comparisons without significantly reducing accuracy. To facilitate practical portfolio selection, we introduce the concept of the efficient region, which identifies subsets of efficient portfolios aligned with specific risk preferences. The algorithm is illustrated empirically using portfolios containing catastrophe (CAT) bonds, which exhibit heavy-tailed and skewed return distributions.

Evgenii Vladimirov (Erasmus University Rotterdam)

Characteristic Function-Based Estimation of Option Pricing Models

In this paper, we develop a novel estimation procedure for parametric option pricing models specified under the risk-neutral measure. We set up our estimation strategy to minimize the distance in a functional sense between the option-implied and model-implied logarithm of conditional characteristic functions, avoiding repeated costly option evaluations. Within a broad class of option pricing models, for which the characteristic function is marginally affine, the model's latent state vector can be concentrated out in closed form. As a result, our estimation procedure is computationally fast and easy to implement, while at the same time exploiting all distributional information contained in an option panel about the risk-neutral dynamics of the underlying asset price. We establish the asymptotic properties of the parameter and state estimators and investigate the finite-sample performance of our method in extensive Monte Carlo simulations. In an empirical application, we illustrate the usefulness of our estimation procedure based on by far the largest option panel, both in the time-series and cross-sectional dimensions, considered in the related literature.

Fenghui Yu (TU Delft)

Explicit signal-adaptive optimal execution quoting strategies

This talk presents optimal execution strategies for sequentially placing limit orders at adaptive quote prices in a limit order book. Unlike classical approaches that focus on trading speed, we study signal-adaptive quoting strategies that account for execution risk and price impact. We consider four settings: risk-neutral execution, execution with running inventory risk, exponential utility, and their combination. For general price impact and inventory risk functions, we show that all problems reduce to a unified structure admitting fully explicit solutions. We illustrate the effectiveness of the optimal strategies using signals learned from limit order book data.

7 Poster presentations

Wouter Andringa (University of Amsterdam)

Semimartingality of signatures and applications to optimal control

We consider utility indifference pricing and hedging in the signature volatility model. In this setting, the Markovian state space contains the signature of the Brownian motion driving the volatility. The derivation of the HJB equation demands semimartingality of this state space. We show that the signature of the Brownian motion is a semimartingale and that the resulting infinite-dimensional HJB PDE has a viscosity solution. We also give some extensions of the semimartingality result by generalising a semimartingale convergence statement to Hilbert spaces. As a future work, we want to numerically solve the high-dimensional PDE by using the Deep 2BSDE method.

Konstantinos Chatziandreou (University of Amsterdam)

Semi-static variance optimal hedging of multi-asset derivatives under affine stochastic covariance models

We develop semi-static variance-optimal hedging strategies for multi-asset contingent claims in multivariate stochastic covariance models. These strategies combine continuous trading in the underlying assets with static positions in a set of auxiliary instruments. Using a multivariate Galtchouk-Kunita–Watanabe (GKW) decomposition, we demonstrate that the semi-static variance-optimal hedging problem naturally decouples into an inner dynamic quadratic hedging problem and an outer optimization. For affine and linear–quadratic stochastic covariance models, we obtain semi-closed-form expressions for optimal hedge ratios and prices. Applications to a range of exotic multi-asset derivatives illustrate the tractability and interpretability of the proposed framework.

Serena Della Corte (Leiden University)

A comparison principle based on couplings of partial integro-differential operators

We prove a comparison principle for viscosity solutions of Hamilton–Jacobi, Hamilton–Jacobi–Bellman, and Hamilton–Jacobi–Isaacs equations with second-order partial integro-differential operators, including generators of (inhomogeneous) Lévy processes. The method combines a test-function version of the Crandall–Ishii lemma with coupling-based estimates to control the difference of Hamiltonians, together with Lyapunov techniques for global behavior. As a consequence, we obtain uniqueness of viscosity solutions arising in stochastic control and differential games with jumps, as in pricing problems under model uncertainty.

Truong Ngoc Nguyen (Utrecht University)

Single- and Multi-Level Fourier-RQMC Methods for Multivariate Shortfall Risk

Multivariate shortfall risk measures yield systemic-risk-aware capital allocations prior to aggregation, but existing Monte Carlo estimators can be computationally expensive. We develop a new class of single- and multilevel numerical algorithms for estimating multivariate shortfall risk and the associated optimal allocations, based on a combination of Fourier inversion techniques and randomized quasi-Monte Carlo (RQMC) sampling, which exploits the geometric convergence of the deterministic optimizer. We establish a rigorous mathematical framework for the resulting Fourier-RQMC estimators, including convergence analysis and computational complexity bounds. Numerical experiments demonstrate that the proposed Fourier-RQMC methods outperform sample average approximation (SAA) and stochastic optimization benchmarks (SA) in terms of accuracy and computational cost across a range of models for the risk factors and loss structures.